A Robotic Swarm for Spill Finding and Perimeter Formation

David J. Bruemmer
Donald D. Dudenhoeffer
Mark D. McKay
Matthew O. Anderson

The Human-System Simulation Laboratory
Idaho National Engineering and Environmental Laboratory
P.O. Box 1625
Idaho Falls, ID 83415-3779, U.S.A.
bruedj,dudded,mark,matthew@inel.gov

Abstract

This paper addresses issues surrounding deployment and tasking of a real-world collective of cost-effective, small mobile robots. To escape the limitations of centralized control, this project distributes control using an innovative, multi-modal communication architecture including acoustical chirping, infrared, and radio frequency transmissions. This paper reports on the use of social potential fields – attractive and repulsive fields emitted by each robot -- as a means to coordinate group behavior and promote the emergence of swarm intelligence as seen in a colony of ants or swarm of bees. A suite of C2 tools, AgentTools, has been developed to enable an operator to inject high-level domain knowledge and guidance into the behavior of the otherwise autonomous robots. The resulting system permits the user to interact with functional groups, rather than issuing commands to each individual robot. Using the real-world robot collective and C2 system, the Idaho National Engineering and Environmental Laboratory has performed experiments to empirically analyze the benefits and limitations associated with the use of many small-scale robots. Experimental results point to fundamental advantages of distributed systems and indicate that our real-world implementation of social potential fields scales well to varying numbers of robots and improves performance in terms of time and reliability.

I. INTRODUCTION

In the near future, it may be possible to produce and deploy large numbers of inexpensive, disposable, meso-scale robots. Although limited in individual capability, such robots deployed in large numbers would represent a strong cumulative force as with a colony of ants or swarm of bees. However, the problem of creating coordinated social behavior from simple, reactive behavior sets is not easily solved. One means that insect societies use to impose order and structure onto the otherwise erratic behavior of individuals is group formation behavior where a spatial relationship is maintained implicitly between adjacent entities as in a flock of birds, a school of fish, or a swarm of gnats.[1] Likewise, we have found that social potential fields provide a means to control a variety of emergent swarm effects including swarm size, swarm density, swarm translation, and the propensity of the swarm to explore new ground.[2, 3] Our work with a collective of 12 small robots shows that social potential fields, although wrought entirely through local interactions and reactive behaviors, can provide a means for coordination and control of a collective as it performs searches in various environments. By modulating these fields through online adaptation or in response to high-level user commands, it is possible to spur dramatic performance improvements in the behavior of the collective.

From an operational perspective, the near-term goal of our work is to develop a team of small disposable robots to assist a human operator in the remote characterization of hazardous or unknown environments. Within this context, small scale distributed robots can reduce cost, remove workers from the dangers of radioactive, explosive, toxic and other hazardous materials, and increase productivity. We predict that multi-robot systems will one day be used across the Department of Energy complex to map and characterize buried waste and retired facilities; to perform routine inspection of critical components; to perform environmental monitoring and building surveillance and to provide rapid-response capabilities in the event of a hazardous spill or radiation leak. The work outlined in this paper takes initial steps towards this vision.

II. RESEARCH ISSUES

A. Previous Work

The goal of using autonomous robots to perform remote characterizations is not unique. Current robotic systems used for this application tend to be highly sophisticated, expensive platforms – typically large to mid-sized robots deployed as single units or in small groups. Existing systems are domain-centric and often require prior instrumentation and/or teams of engineers to operate and service. They are expensive to manufacture, transport, and operate, and, consequently, are undesirable for rapid response in remote characterization tasks where the robots often cannot be recovered because of exposure to hazards. One innovative approach to the remote characterization problem was developed through a partnership between Oak Ridge National Laboratory (ORNL) and the Idaho National Engineering and Environmental Laboratory (INEEL). This effort produced MACS (Mobile Automated Characterization System) and RACS (Reduce Access Characterization System).[4] The large robot, MACS, explores and characterizes the building, deploying the smaller RACS robot into rooms and areas where the larger robot has limited mobility/access.

While the MACS and RACS team offers many benefits over strategies that employ only one system, such approaches do not exploit the benefits of fully distributed systems:

- *Emergent Behavior* – As in a colony of ants, intelligent, complex behavior emerges from the interactions of multiple robots each driven by simple behaviors.
- *High Fault Tolerance* – By distributing the task across a loosely coupled population of robots, the collective can succeed even when particular robots fail.
- *Redundancy* – The behavior of each robot can be validated / duplicated by its peers.
- *Cooperative Behavior* – We can exploit synergistic behavior impossible with only one to several robots.
- *Modulated Diversity* – As in biological systems, an appropriate level of diversity adds richness to the capabilities of the collective and makes it more robust to environmental changes.
- *Low Cost* – Small scale robots can potentially be used as a disposable resource

Swarm intelligence offers a means to achieve these benefits. The expression 'swarm intelligence' was first used by Beni to describe systems where many simple agents generate patterns and self-organize through nearest neighbor interactions.[5] More recently, Bonabeau, Dorigo and Theraulaz have supplied a useful definition of the term as "any attempt to design algorithms or distributed problem-solving devices inspired by the collective behavior of social insect colonies and other animal societies."[6] Borrowing on previous work by Tsetlin we can identify four main characteristics of swarm behavior including randomness, decentralization, indirect interaction, and self-organization.[7] Throughout this paper, we discuss how each of these ingredients impact performance of our robot collective.

B. A Basis in Biology

Insects do not rely on sophisticated internal states, directed communication, global position information, or range information – the hallmarks of most state-of-the-art robotics efforts.[1] Likewise, our approach has been to abandon tools such as wheel encoders, laser range finders, sonar, vision, GPS, compass, etc., which have proven useful for mid-sized to large robots and embrace an entirely different paradigm that relies on simple, local, interactions displaced onto the environment rather than internal computation. For our robots, these interactions take place through entomologically inspired modes of perception and communication including chirping, detection of other robots' shadows, antennae-like touch sensing and moisture detection. Within this embodied approach, the robots learn to respond appropriately to fluctuations in sound and light; in fact, obstacle avoidance and a variety of social behaviors including searching, spill convergence, and perimeter formation are all dependent on the robot's ability to both recognize and instigate these fluctuations.

Insect behavior is robust to environmental changes because insects exhibit a tight coupling between perception and action. Very little processing occurs from when a cockroach perceives a sudden change in lighting and when it moves. Likewise, our robots do not make deliberative, high-level decisions about the task, but react to their environment using fast, responsive behaviors that are domain independent and robust because they do not rely on sophisticated internal processing. We believe our approach is especially useful for rapid response type missions where little is known about the environment and/or there is insufficient time to custom tailor a robotic solution.

Figure 1 below shows part of the robot collective investigating the floor of a DOE lab facility.

Fig. 1: Robots swarm as they explore a DOE building.

III. IMPLEMENTATION

A. Robot Hardware

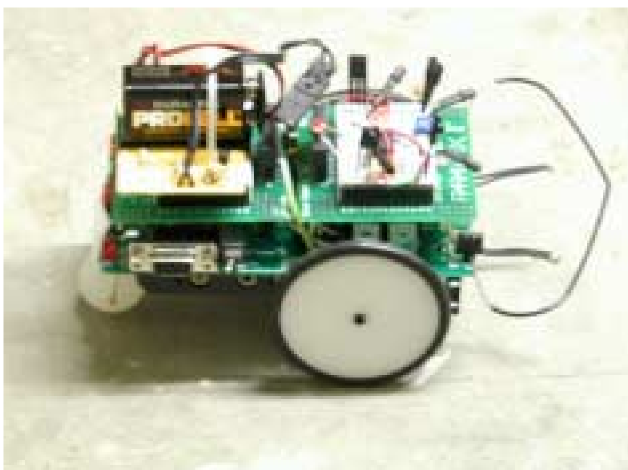

Fig. 2: Instrumented GrowBot

The platform chosen for this research project is based on the basic stamp by Parallax. Our current research platform has two of these processors: one for communications and the other for navigation. In addition, the robot is equipped with a spill detection sensor, two bump sensors, two whisker-like light sensors, four IR sensors for obstacle avoidance, a ring of IR for local communications, a piezoelectric speaker and two directional hearing aid microphones, one in the front and one in the rear. Also, the robot collective includes "sergeant" robots that are specialized for communication. Radio frequency transmission capabilities allow the sergeants to receive commands from a human operator. The sergeants then use IR transmission to disseminate these commands to the "privates." To facilitate user interaction with large numbers of robots, privates can be assigned to swarm around a particular sergeant as it moves to accomplish search and detection tasks.

B. Social Potential Fields

The INEEL has implemented social potential fields on a collection of 12 robots using a combination of IR obstacle avoidance, light sensing and audible chirping. Through these behaviors, each robot exerts both an attractive and repulsive force field. The attractive field, based primarily on sound, can either discourage robots from moving too far away (an essential aspect of stable swarming behavior) or can actively pull other robots towards itself through a "come hither" chirp emitted by a robot that has found an area of interest, such as a spill. The repulsive field discourages robots from coming too close and is based on sound (robots avoid chirps above a certain volume) and the various obstacle avoidance sensors, which include infrared, light sensing, and bump sensing.

Online Learning

The elusive goal was to develop behavior that could implicitly (i.e. without a map, internal representation, directed communication, or centralized control):

- Minimize redundancy and interference
- Maintain a beneficial level of social interaction
- Adjust each robot's willingness to explore
- Automatically adapt individual robots to different environments and varying numbers of robots

To accomplish these aims, we developed a form of online adaptation that provides the swarm with a means to automatically regulate itself. Positive and negative feedback is supplied to each robot by an internal critic, invoked at regular time intervals in order to continually adjust sensitivity to light and sound fluctuations. If the robot is too sensitive to these fluctuations, it appears "timid" and will fail to cover new ground. On the other hand, if the robot becomes unresponsive to such fluctuations, it will not effectively avoid collisions with obstacles other robots. Perception of real world light intensity and sound fluctuations offers a perfect means to modulate levels of randomness and diversity – key components of swarm behavior – into the robots' behavior. [5,6] By adjusting the level of randomness, the online learning system can modulate certain emergent properties of the swarm such as swarm density, swarm translation, and swarm convergence. It also adapts the swarm to new environments and promotes full coverage even in obstacle rich environments.

Fig. 3: The Parent robot deploys the smaller robots through a doorway into a large DOE facility.

Fig. 4: The view from the Parent robot. The operator uses this feedback to guide the sergeant through the door as the privates swarm behind it.

C. Command and Control

Although self-regulation offers dividends in terms of robustness, the resulting diversity is not easy to predict or precisely control. Imagine trying to develop command and control for a colony of ants or a swarm of gnats. Unlike the insect world, the robotic system must interact cooperatively with human operators. Ideally, the user should not be required to task individuals, but should be able to abstract group command and control functions. To support this need for high-level tasking, INEEL has developed AgentCDR, a hierarchical command and control tool that includes human-centric visualization tools, iconographic representation of robots, GUI controlled group assignment, operation planning tools, and system status alerts for communication failure. Furthermore, the privates are not dependent on the sergeants or on the human operator for continuous communication and can function autonomously in the absence of sergeant or user input. This flexibility supports mixed-initiative control and allows AgentCDR to balance the needs and limitations of the robots, C2 structure, and the human operator(s).

D. Parent Robot

One of the issues in utilizing small robots is control of their initial placement within the environment. To address this deployment problem, the INEEL has developed a Parent robot – a much larger and more capable robot -- that can deploy the robots by emitting a “follow me” chirp. In turn the smaller robots utilize a combination of an IR-based follow behavior and a “chirp follow” behavior to track the Parent. The Parent robot deploys the robots into a building, and then assumes a monitoring mode. Using an autonomous tracking behavior, the Parent provides visual feedback on a particular group by following a certain distance behind a specified sergeant. In terms of the operational scenario, the ability to autonomously provide visual feedback is a crucial form of support for the operator using AgentCDR. In a recent demonstration of the system in a DOE building, the ability of the Parent to provide feedback on swarm behavior allowed the user to accomplish difficult tasks such as guiding a group of robots through a door and into a new area of the building. Rather than adding complexity to the task, the Parent robot alleviates cognitive load for the operator and augmented overall swarm utility.

IV. EXPERIMENT

Before we could fully understand the effects of social potential fields, we needed some way to empirically measure the performance of our swarm as it accomplished search and perimeter formation tasks. Our goal was to show that the social interactions wrought through our implementation of social potential fields could produce desirable emergent effects that helped the robots accomplish their task. We needed some means to quantify the performance benefits achieved though emergent social effects and chose to do so in terms of a coverage efficiency experiment. The experiment focused on the autonomous searching behavior of the privates and did not examine command and control issues related to infusing the system with human knowledge and guidance.

A. Methods

The primary challenge was how to acquire empirical, objective data on the behavior of the robots. To meet this need, we constructed an environment consisting of an eight by eight foot walled enclosure with a floor covering consisting of large sheets of white-board. Each robot was instrumented with a Velcro sponge pad, which allowed us to securely attach a dry erase marker to the rear of the robot. Each robot was fitted with a different color marker to differentiate its path from the others. The marker provided an effective means to capture “ground-truth” on the movements of each robot and the cumulative effect on the resulting area coverage. Figure 5 illustrates the test bed environment.

Fig. 5: A single robot explores the test bed environment

To complete the coverage task, the robot(s) were required to fully explore the floor of the test bed described above. We considered several means of ascertaining coverage and decided that full-coverage would be defined as “no unmarked space remaining into which a robot could fit lengthwise.” Throughout the experiment, four rectangular, cardboard obstacles of varying sizes remained fixed in position within the testbed.

We ran five trials with one, two, three, four, six and nine robots. For each trial the robot(s) were placed in the same corner and were all started within a few seconds of each other. For each trial, we recorded the total time required to achieve full coverage and then wiped the test-bed clean.

B. Results

TABLE I
TIME REQUIRED FOR TASK COMPLETION (MIN)

Robots	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
1	77	53	110	113	170
2	33.78	56	50.93	42.33	32.67
3	21.93	24.63	18.47	14.93	18.33
4	14.9	12.63	10.25	14.76	19.0
6	6.53	6.46	6.78	7.58	13.20
9	11.75	6.36	8.02	5.47	7.63

Table 1 provides the time required for each run.

Figure 6 shows that the average time required to achieve full coverage decreases drastically as the number of robots increases.

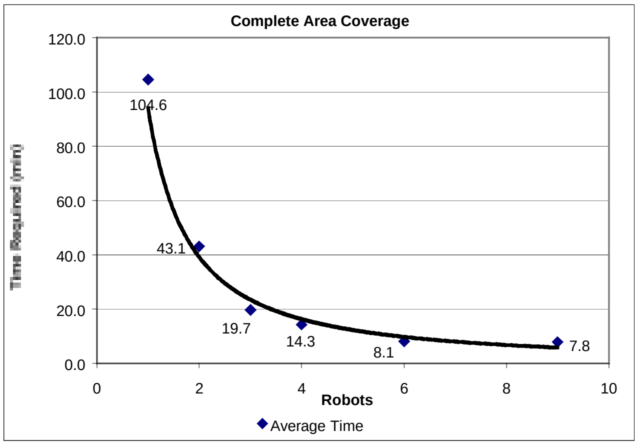

Fig. 6: Average time for complete coverage

If we define performance as the reciprocal of time required, we can present the data in terms of overall performance / the number of robots. Figure 7 shows the system performance for each trial divided by the number of robots used. This “performance per robot,” metric is commonly used to discuss how the synergistic effects of cooperation changes with varying numbers of robots.[8]

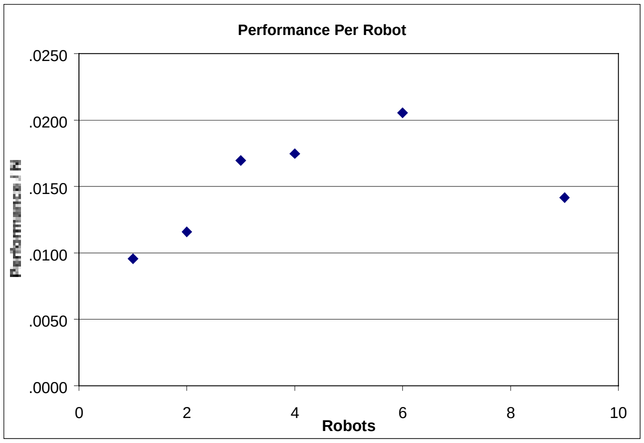

Fig. 7: Overall performance for each set of trials / n robots.

C. Discussion

While a 64 sq. ft. environment may seem small in comparison to many operational environments, complete coverage proved to be a stiff requirement. The time required for a single robot to achieve full coverage varied drastically from trial to trial. Indeed, the distribution range between trials may be, in and of itself, a significant result. Besides reducing the overall time required to search an area, the use of multiple robots renders overall task performance more reliable. Figure 8 shows that the use of multiple robot improves reliability of the system.

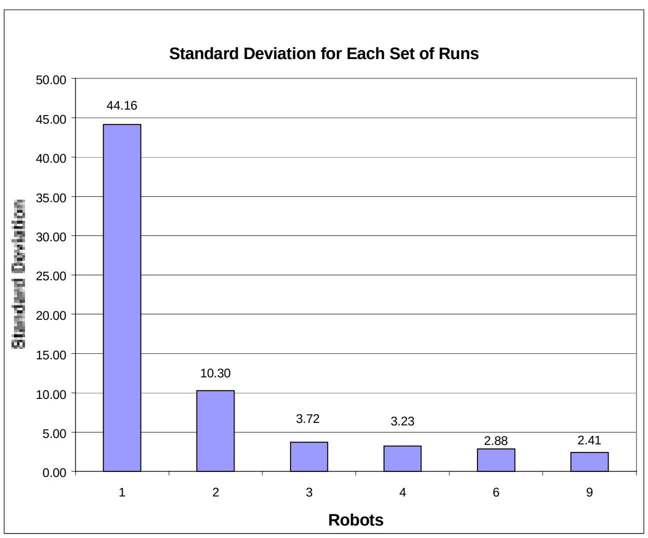

Fig. 8: Standard deviation for each set of trials

Our data show that the performance per robot increases as we add more robots, indicating that there is a "synergistic" effect emerging. This indicates that we are indeed benefiting from the social effects of multiple robot interaction and that these effects grow as we add more robots. However, our results also indicate that the benefit to adding additional robots extends only to a point after which the synergistic effects begin to be offset by the detrimental effects of increasing interference. Figure 7 shows that the performance per robot augments through six robots but then begins to diminish. By nine robots, the performance per robot has drastically decreased.

One aspect of the experimental design, which had a significant impact on the results, was our decision to measure the time needed for complete coverage as opposed to measuring percent of coverage after a fixed time had elapsed. At the beginning of a coverage task, much new ground is being covered, while by the end, the robot(s) are covering very little new ground. Thus, the difference between 1 and n robots appears most stark when we wait for that last little bit of ground to be covered. With 1 robot, the last remaining small area could take as much time to cover as the whole rest of the environment. With multiple robots this effect is greatly mitigated.

These results suggest that use of multiple robots can be a great advantage for search and detection tasks. However, before we can draw definitive conclusions regarding these speculations, it is necessary to reproduce the experiments with larger numbers of robots in different environments. Reproducing this experiment will allow us to ascertain which results generalize across environments and which are a function of the specific study reported here.

V. CONCLUSION

In a variety of environments, the INEEL has effectively demonstrated that an operator can use the multi-robot system, including the Parent robot, Sergeants and Privates, to search through multiple rooms and converge upon a mock spill. At the 2001 American Association for Artificial Intelligence (AAAI) Conference in Seattle, WA, the INEEL demonstrated the ability of a team of robots to autonomously locate and form a perimeter around a water spill within an Urban Search and Rescue test bed designed and built by the National Institute of Standards and Technology (NIST).[9,10] After release, the robots began to disperse using their social potential fields to implicitly divide the environment. Once one robot found the spill and began to emit an audible signal, congregation and perimeter formation occurred quickly as the other robots were drawn to the "come hither" chirp.[11] Figure 9 below shows the result of a similar demonstration performed within a large, cluttered building at the INEEL.

Figure 9: Robots deployed in a cluttered DOE facility form a perimeter around a water spill.

In the future, the INEEL will investigate how to transition these capabilities towards specific Department of Energy applications. A small radiation sensor has recently been interfaced to the Privates, which will allow them to map radiation gradients and converge on the source. In addition, work is underway to move the processors, sensors, communication and existing behaviors onto a more capable platform, based on an inexpensive, off-the-shelf toy monster truck chassis. These new robots will be fast (up to 40mph), rugged and capable of handling rough terrain. Our plan is to deploy a team of 100 robots within the next 18 months.

ACKNOWLEDGEMENTS

This work is supported through the INEEL Long-Term Research Initiative Program under DOE Idaho Operations Office Contract DE-AC07-99ID13727 and through DARPA's Software for Distributed Robotics project under research contract J933.

REFERENCES

1. Wilson, E. O. 2000. *Sociobiology: The New Synthesis.* Cambridge, MA: Belknap Press.

2. Reif, John H. and Hohgyan Wang. 1999. Social potential fields: A distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems* 27: 171-194.

3. Dudenhoeffer, Donald D. and Michael P. Jones. 2000. "A Formation Behavior for Large-Scale Micro-Robot Force Deployment". In *Proceeding of WSC'00*.

4. Anderson, M.O., M.D. McKay, and B.S. Richardson 1996, *Multirobot Automated Indoor Floor Characterization Team. Proceedings of the 1996 IEEE International Conference on Robotics and Automation, pp. 1750-3.*

5. Beni, G. 1988, "The Concept of Cellular Robotic Systems." In *Proceedings 1988 IEEE Int. Symp. On Intelligent Control*, 57-62. Los Alamitos, CA: IEEE Computer Society Press.

6. Bonabeau, Eric, Dorigo, Marco, Theraulaz, Guy. 1999. *Swarm Intelligence,* New York, NY: Oxford University Press

7. Tsetlin, M. L. *Automaton Theory and Modeling of Biological Systems.* New York, NY: Academic Press, 1973.

8. Balch, T. "Hierarchic Social Entropy and Behavior Differences: New Measures of Robot Group Diversity." In *Proceedings of Performance Metrics for Intelligent Systems Workshop,* Gaithersburg, August 2000.

9. Murphy, R., Casper, J., Micire, M., Hyams, J. "Assessment of the MIST Standard Test Bed for Urban Search and Rescue." In Proceedings of *Performance Metrics for Intelligent Systems Workshop*, Gaithersburg, 2000.

10. Jacoff, A., Messina, E., Evans, J., "Astandard Test Course for Urban Search and Rescue." In Proceedings of *Performance Metrics for Intelligent Systems Workshop*, Gaithersburg, 2000.

11. Bruemmer, D. J., Dudenhoeffer, D. D., Marble, J. "Spill-Finding and Perimeter Formation in a Distributed Team of Small Robots." In *2001 AAAI Mobile Robot Competition and Exhibition Workshop*, Seattle, August 2001.